The topology of mental states.
Combining data science and graph theory to reveal neural networks for cognitive functions.

Andrea Insabato

Universitat Pompeu Fabra
Theoretical and Computational Neuroscience
Center for Brain and Cognition
Roc Boronat, 138
08018 Barcelona, Spain,

Columbia University
Italian Academy for Advanced Studies
Center for Theoretical Neuroscience
1161 Amsterdam Ave.
New York NY 10027, USA

October 30, 2017

Abstract

In this article we give a general introduction to the study of cognitive functions and dysfunctions from the perspective of data science and network science.

We will use functional magnetic resonance imaging data since it is a privileged method to observe whole brain activity in a non-invasive way.

The analysis of specific brain networks will provide a quantitative method to investigate the neural correlates of cognition and at the same time can be used to develop new diagnostic and therapeutic procedures for individual neuropsychiatric patients.

The extraction of specific networks has to face the problem of a mixture of variability in brain activity. While extracting the network, for example, related to a given pathology, we want to assess the variability in this network between individual patients and filter out the variability due to repeated sampling. In order to face this problem we propose to simultaneously classify subjects and conditions (behavioral or pathologic).

We base our classification pipeline on a network model of brain activity that enable us to extract the effective connectivity between brain regions. In contrast to recent attempts to classify individuals based on functional connectivity, our approach provides more stable and reliable estimation of connectivity to the purpose of the classification.

Here we show that our method is able to classify a large number of subjects using few recording sessions and at the same time the behavioral condition of the sessions (resting or movie viewing). In addition we extract the networks that underlie the two classifications: The subject network is almost fully connected with several central nodes, located in frontal and cingulate regions; in comparison, the condition network is segregated into small isolated components.

Finally we discuss some refinements of our pipeline that are needed for small sampling size cases.

Contents

1 Introduction

This article is meant to be a general introduction to the study of cognitive functions and dysfunctions from the perspective of data science and network science. While the idea of the brain as the physical substrate of *mental* activity is quite old and more or less well established and the idea of the brain as a network of interconnected units is, although not as old, probably more widespread and shared, the study of whole brain networks underlying cognitive functions and brain disorders is only becoming possible very recently thanks to the acquisition of huge amount of data (often referred to as "big data") and to the development of adequate quantitative methods to analyze such data. The specific focus adopted here will be on estimating connectivity from functional magnetic resonance imaging (fMRI) data as a privileged method to observe whole brain activity in a non-invasive way. However the extension to other type of data, such as calcium imaging, is straightforward.

Relevance. At least two research areas make our study of the utmost importance. On one side it will provide a quantitative method to investigate the neural correlates of cognition at the level of whole brain networks, hence advancing our understanding of cognitive functions. On the other, it will foster the development of individualized medicine for neuropsychiatric diseases using imaging techniques, an idea that has been recently put forward [1].

State-of-the-art. Blood-oxygen-level dependent (BOLD) signals in fMRI have been used for more than two decades to observe human brain activity and relate it to functions [2]. Even at rest, the brain exhibits patterns of correlated activity between distant areas [3]. The functional connectivity (FC) measures the statistical dependencies between the BOLD activities of brain regions, which has then been studied for subjects performing tasks and compared with the resting state. Following fundamental discoveries about brain functions, fMRI has been increasingly used to complement clinical diagnostic for neuropathologies [4]. Resting-state fMRI has also been found to be informative about neuropsychiatric disorders [5]: alterations in FC correlate with and can predict the clinical scores of several diseases [6, 7].

Recent studies have focused on the reliability of these FC measures recorded from the same subject over successive sessions [8, 9]. Consistent differences between subjects (with individual stability) allow subject identification using recorded FC as a fingerprint [10]. Moreover, this subject specificity may even be enhanced in task-evoked activity [11]. A recent prospective study about the evolution of psychiatric disorders emphasized individual specificities in the FC stabilization during childhood (irrespective of the disease) [12], whereas traditional group-averaging aims to remove the individual differences to obtain task-specific [13] or pathology-specific [14] signatures.

Anatomy of the problem. We should now dissect the general problem prospected above (the study of whole brain networks underlying cognitive dis-functions) into different subproblems. As we have seen, the mixture of session-to-session, subject-specific and condition-related variability in FC is a crucial issue in particular when only a few sessions per subject can be recorded, such as for clinical diagnostic. So the original problem can be translated into distinguishing these different sources of variability. In particular classifying subjects and conditions [1] would be a practical solution to this problem, since

GLOSSARY

fMRI: the activity of neurons is visualized through the induced magnetic field. Current spatial resolution of fMRI is in the order of mm, allowing to study large ensembles of neurons such as hypercolumns.

BOLD: The activity of the brain requires oxygen. Oxygenated blood has different magnetic properties compared to de-oxygenated blood. These differences can be recorded by an fMRI scanner (the so called BOLD signal), thereby providing a proxy for the level of activity of each recorded brain region.

FC: Pattern of statistical association between brain areas. While structural connectivity represents the physical connections (the fibers) between areas, FC reflects the use of those physical connections when the brain is active. While structural connectivity remains the same even when the brain is dead, FC changes depending on the type of activity.

[1] We use the neutral word "condition" to indicate the condition a subject might be in, be it mental, e.g. dreaming, counting, remembering, or a pathological/healthy condition.

classifying both we would be filtering out the remaining session-to-session variability.

Approach towards a solution. As already mentioned here we will take a whole brain approach. Distributed signatures in FC across the whole brain have been observed in memory tasks [15] or when the subject experiences psychological pain [16]. Moreover, the etiology of many mental disorders is unknown: they are suspected to arise from network dysfunction, as reported for large-scale FC alterations in patients with schizophrenia [17]. These examples strongly point in favor of whole-brain approaches to study high-level cognition and brain diseases; in contrast, focusing on a few cortical areas only to test hypotheses may not capture sufficient information and network effects. Such whole-brain approaches typically involve a large number of parameters to estimate, which may impair the robustness. One aim of the present study is to provide a practical answer to this trade-off dilemma.

The idea underlying the study of FC in the broad sense lies in that it reflects how brain areas dynamically bind to exchange and process information [18, 19]. To move beyond a phenomenological description of FC, our method relies on a mathematical network model of BOLD time series [20] and allow to FC to be decomposed into changes in network connectivity, called effective connectivity (EC) and local fluctuating activity. As with FC, a crucial issue for EC is whether the estimated model parameters are reliable across several sessions for the same subject [21], which determines whether they can predict the subjects' identities in practice [22].

We will first develop the method for simultaneous subject and condition identification for simple conditions. In particular we will use a dataset where subjects can be in two conditions each during 4 minutes: resting or viewing a movie. We will then refine the method for its application to diverse conditions such as healthy/pathological states or short duration *mental* states, such as memory retrieval, visual exploration, working memory, decision-making, etc. In particular we expect an increased complexity from short duration states.

Summarizing, the steps towards the accomplishment of our objective are:

1. Reliably extract whole brain EC;
2. Identify individuals;
3. Identify conditions;
4. Extract subnetworks that characterize individuals and conditions;
5. Refine the method for short duration states;

Classification: In machine-learning is the task of predicting the category a new object belongs to, given previous examples of different objects and a categorical scheme. As an example, the spam filter of your email uses a classifier that categorizes incoming email in *spam* and *non-spam*.

EC: Although not a formally defined concept, it refers to the real connections employed by a network during its activity. In practice usually the parameters of a network model are adjusted in order to reproduce experimental data and the connectivity of the model is used as an estimation of EC.

2 A new method for the extraction of networks related to subjective identity and condition

The vast majority of the results presented here have been presented in a recent technical report [23]. Here we will avoid a detailed description of methods (all necessary technical information will be presented together with the results) and

Dataset name	Acquisition	Number of subjects	Sessions per subject	Session duration
Dataset A1	Day2day project	6	40-50	5 minutes
Dataset A2	Day2day project	50	1	5 minutes
Dataset B	CoRR	30	10	10 minutes
Dataset C	[24, 4]	19	3 resting; 2 movie	10 minutes

Table 1: Datasets used in the study. Dataset A (A1+A2) was used to test the robustness of subject identification in session-to-session variability. Dataset B was used to test the generalization capability of the identification procedure for a larger number of subjects. Dataset C was used to extract both individualized and behavioral signatures.

we refer the reader to the technical report for a full description the procedures.

2.1 fMRI data

In this study we used fMRI data from the three datasets, where several subjects underwent multiple sessions of fMRI recording, as described in Table 1. The processing of the data is represented in fig.1A. For each individual session data were corrected for artifacts and parcelated in regions of interest (ROIs). The mean BOLD signal in each ROI was extracted and classical functional connectivity (corrFC) was calculated using the pairwise Pearson correlation coefficient (PCC) between the time courses of BOLD in each ROI. Finally, as show in Fig.1A, we obtain an NN symmetric matrix for each recorded session (N=116 for Datasets A and B, N=66 for Dataset C), where each element at row i and column j of the matrix represents the PCC between the ith ROI and the jth ROI. The matrix is symmetric since PCC is symmetric $(PCC(i,j) = PCC(j,i))$.

PCC: Measures the strength of linear statistical association between two variables. In the case where variable x increases and variable y increases PCC will be high and positive (with a maximum of 1), when x increases and y decreases PCC will be negative (with a minimum of -1). When the variables are not associated PCC will be 0 or very small.

2.2 Network model and effective connectivity extraction

The model is a network where each element, corresponding to each ROI, is is governed by an Ornestein-Uhlenbeck (OU) equation. To the aim of extracting the EC we used the whole-brain dynamic model [25] depicted in Figure 1B. Each ROI receives input from other ROIs according to the connectivity scheme defined by the EC, which is estimated from the data. The model can be simulated and the FC of the model can be compared to that of empirical data. In the model, the global pattern of FC arises from the local variability Σ_i that propagates via the network connections EC_{ij} (from j to i). To fit each fMRI session, all relevant EC_{ij} and Σ_i parameters are iteratively tuned such that the model FC best reproduces the empirical counterpart. A detailed description of the model and the maximum-likelihood estimation procedure is provided in [25].

Mathematical model: A set of differential equations that describe the evolution in time of some variables, possibly coupled, and their relationship to other static variables. Static variables can be determined by empirical measures or unknown, i.e. their value can be freely set.

Simulation: The solution of the differential equations of a model that gives the evolution in time of the variables of interest. Since (at least some) variables represent measurable quantities, the process *simulates* the observable world.

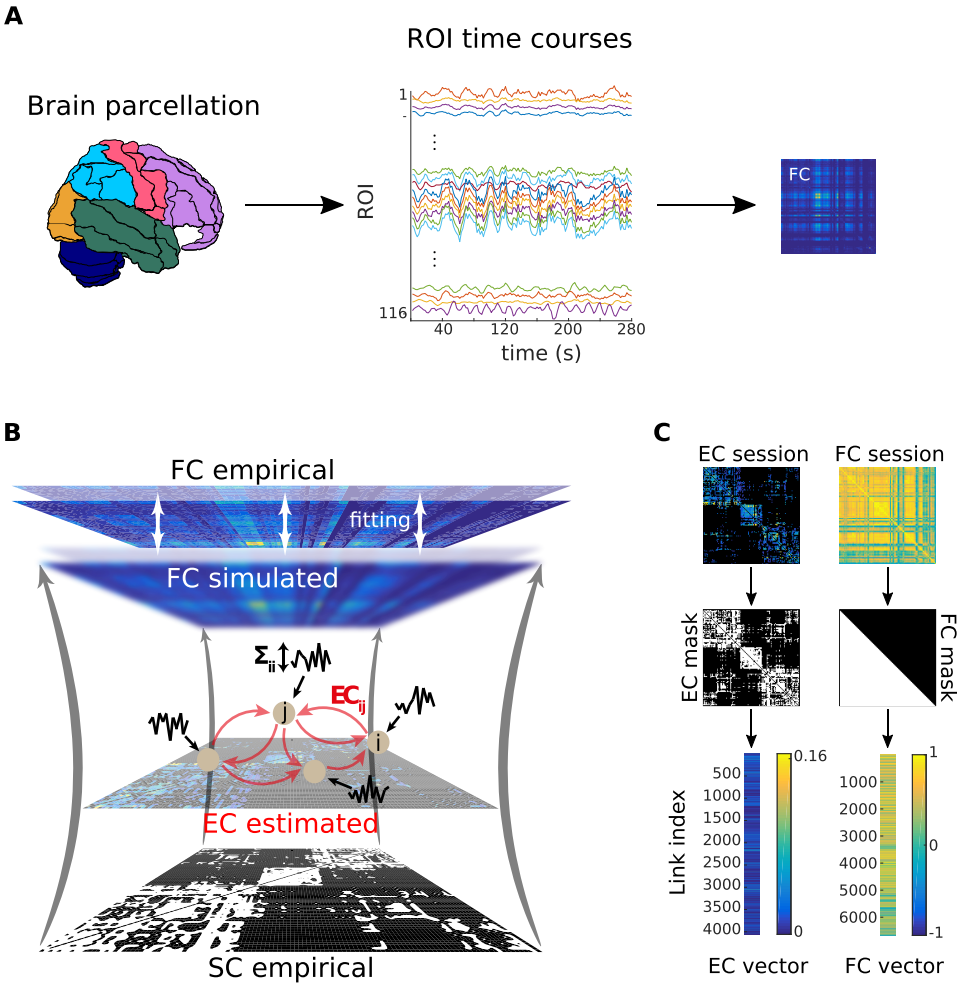

Figure 1: After a standard pre-processing pipeline, a parcellation covering the whole-brain is applied to extract BOLD time series with each color representing an anatomical subsystem of several ROIs. B) Whole-brain network model. The local fluctuating activity propagates via the recurrent EC to generate the correlation patterns at the network level. The fitting procedure iteratively tunes EC and Σ such that the model best reproduces the empirical FC. C) Each corrFC matrix is symmetric and has all diagonal elements equal to 1, so that only 6670 independent links are retained for classification (lower triangle). Likewise, the EC matrix has 4056 non-zero elements that are used in the classification (density of 30%).

2.3 Subject identification

Robust subject identification for ∼100 subjects was pioneered by a recent publication [11], relying on a k-nearest-neighbor (kNN) classifier with k=1 and PCC as metric. In contrast with previous studies using 1NN [11, 21, 12], our method relies on a multinomial logistic regression (MLR) classifier, a classical tool in machine learning. MLR uses a linear model to predict the probability that an input sample belongs to a class (subject here).

In classification algorithms the problem of overfitting describes the situation where the algorithm performs very well with the data it is trained with, but fails to generalize to new samples. Due to the high dimensionality of the connectivity measures, it is essential to control for overfitting with an appropriate training and test procedure. Our train-test procedure and the use of large test-retest datasets unlike previous studies [11, 21, 26] aims to provide a trustworthy characterization of the quality of the classifiers. Figure 2A describes the train-test procedure for the identification of subjects: 1) fMRI sessions (EC in the figure) are randomly split in training and test datasets; 2) after preprocessing (orange arrows) involving within-session z-score followed or not by PCA, the classifier is optimized as illustrated for the MLR with boundaries that best predict the training dataset; 3) test set is used to verify the generalization capability of the classifier (blue arrows), by measuring to which extent the classifier boundaries, estimated with the train set, correctly classify single sessions from the test set.

We first used Dataset A1 and increased the number of training sessions per subject from 1 to 40 to evaluate how many training sessions are necessary for satisfactory accuracy. As shown in Figure 3B, EC (in red) outperformed corrFC (in blue) by more than one standard deviation (shaded area around the curve), for both MLR and 1NN. Moreover, almost perfect classification was reached with MLR for only 5 training sessions, whereas 10-15 were necessary for 1NN. This is important when only a few training sessions per subject are available, as expected with clinical applications. Figure 3C displays the classification accuracy for Dataset B, used to verify the robustness with respect to the number of subjects to be classified. We trained the classifiers with 1 session per subject and evaluated the performance varying the number of subjects from 2 to 30 (test set comprised the remaining 9 sessions per subject). Again, EC is more robust than corrFC: while performance with corrFC rapidly deteriorates as the number of subjects is increased, classification using EC is barely affected by the number of subjects. This is our core technical result: EC and MLR largely outperform corrFC and 1NN, respectively.

2.4 Network of subject identification

An important advantage of the MLR over kNN is its efficiency in characterizing the links that contribute to the classification. We used recursive feature elimination (RFE) to rank the links according to their weight in the classification and then chose the lowest number of links that achieved the maximum classification performance. The resulting support network for dataset A1 had 18 links, compared to 44 links for dataset B. In both cases, subject identification using only those links achieved perfect accuracy. The two support networks are shown in Figure 3A in the same matrix: remarkably, the networks are very sparse

Fitting: Or "parameter estimation", is the adjustment of free variables of a model in order to make the simulation as close to the empirical data as possible.

Inference: When a model is fitted to experimental data (i.e. it can reproduce them to a certain extent), the value of free parameters can be inferred, e.g. in a network model the connectivity is typically not measurable but can be inferred once the model has been fitted.

OU: Is a mathematical model that describes the velocity of a Brownian motion, i.e. the motion of a random moving particle. In practice, the dynamic of the process tends toward an equilibrium but is otherwise noisy given the random fluctuations of its input. When multiple OU processes are connected complex dynamics can emerge.

z-score: Transformation of a distribution that allows to highlight the fluctuations of data around the mean.

PCA: Technique that allows to reduce the dimensionality of a dataset by rotating and projecting the data in a new space.

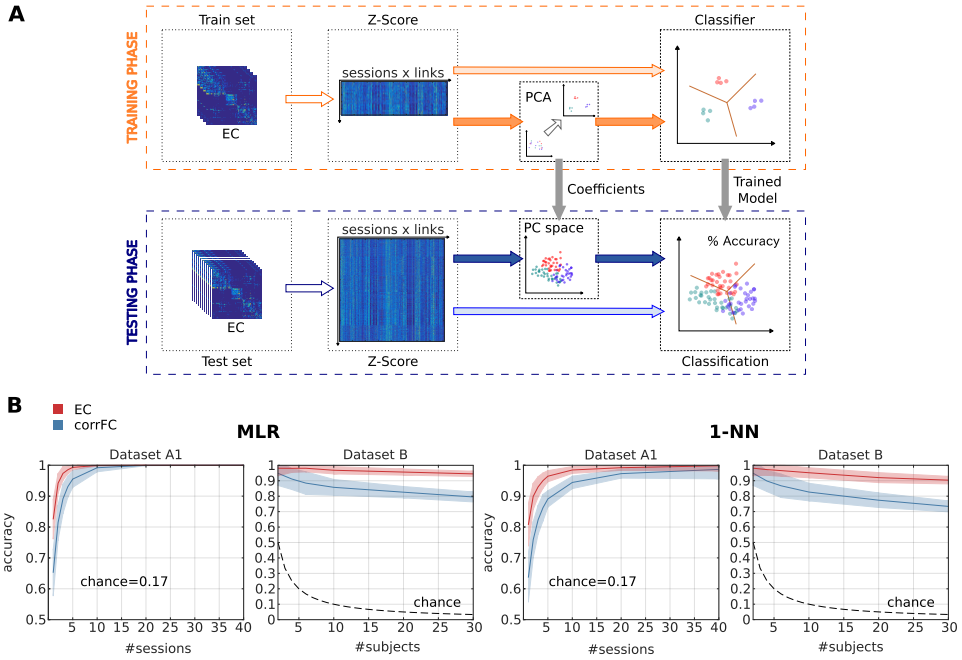

Figure 2: Subject identification using EC and FC. A) Classification pipeline used to assess the generalization of performance. The full set of connectivity measures (here EC) over all fMRI sessions was split into two groups: a train set and a test set. We use z-scores calculated over the elements of each session matrix. We trained the classifier with or without previously applying PCA and evaluated the classification accuracy on the test set. B) Performance of multinomial logistic regression (MLR, left panel) and 1-nearest-neighbor (1NN, right panel) classifiers when increasing the number of sessions per subject used as training set with Dataset A1. The mean (solid curve) and standard deviation (colored area) were calculated for 100 repetitions with cross-validation. C) Same as B when varying the number of subjects using Dataset B, using a single training session per subject (leaving 9 sessions per subject as test test).

and non-uniformly distributed across the whole brain. This is the signature of the most subject-discriminative ROIs: frontal and cingulate cortices, as well as the temporal and occipital regions, seem to play a major role here. It is worth noting that the adjacency matrix is not symmetric, which implies different roles for nodes as receivers (especially frontal ROIs) or senders (cingulate).

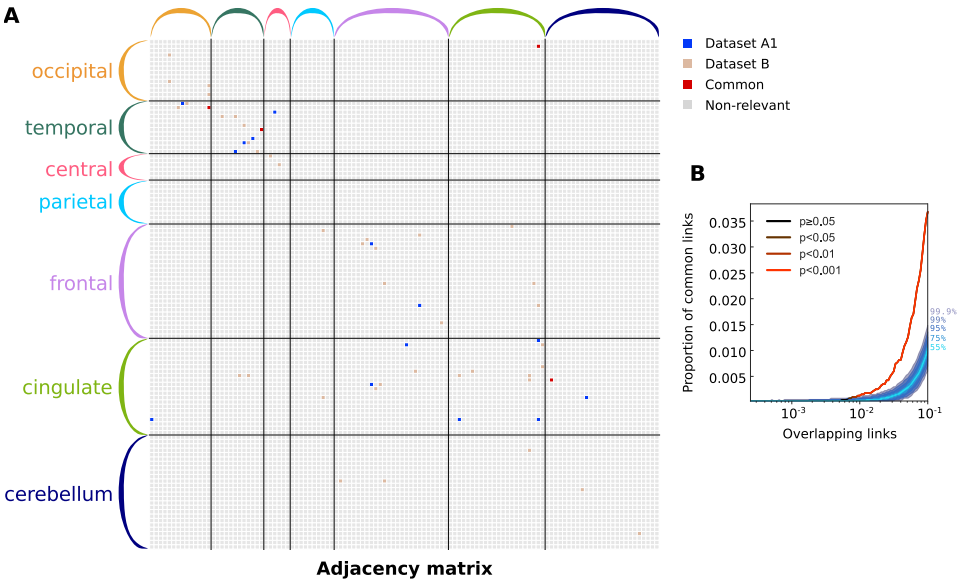

Figure 3: Networks that support subject identification. A) Extracted links that contribute to the classification with both datasets, obtained using recursive feature elimination (RFE). The ROIs are grouped in anatomical pools. B) Overlap between the two signatures for Datasets A1 and B as a function of selected links. The curve represents the amount of common links in the data. Shaded areas represent different quantiles of the surrogate distribution of common links under the null-hypothesis of random rankings. The color of the curve indicates the probability of the corresponding amount of common links under the null-hypothesis (here p-value ¡ 0.001 when considering more than 1% of the total links, namely 40 links).

Network: A set of nodes connected by links. As shown below, network can be represented with dots and arrows or in matrix form, where each element of the matrix correspond to one link and its value is 1 if the link exists or 0 if it doesn't.

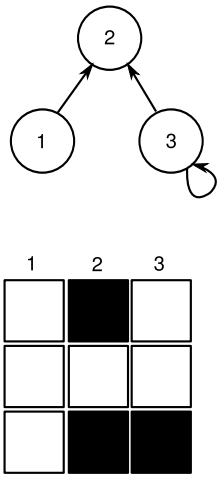

The sparsity of the signature in Figure 3A hides the fact that the rankings for Datasets A1 and B are close (PCC=0.59, p-value¡¡10-50), indicating that similar neural networks characterize individuals in two disjoint sets of subjects; see also Figure 4 that illustrates the correspondence at the level of anatomical groups. To further measure the overlap between these networks, we selected the subset of links with the highest ranking for each dataset and computed the number of common links. Figure 3E shows that the proportion of common links exceeds by far its expectation under the hypothesis of random rankings (shaded gray area). This indicates a good agreement between the support networks from the two datasets even at the single-link level.

Hypothesis testing: To test a hypothesis, the opposite hypothesis, called null-hypothesis, is formulated, encompassing all possible cases but the one under test. The probability of a real measurement under the null-hypothesis is calculated and if it is very small (usually < 0.05) the null-hypothesis is rejected.

2.5 Condition identification and related network

Finally, we used Dataset C to extract a signature for the subject identity and another for the behavioral condition. This is schematically depicted in Figure 5A, with three fictive dimensions: the information about subject identity corresponds to the x-axis and information about the condition to the z-axis; the session-to-session variability, that should be ignored, spreads along the y-axis. In this idealized low-dimensional scenario, it is possible to classify a session

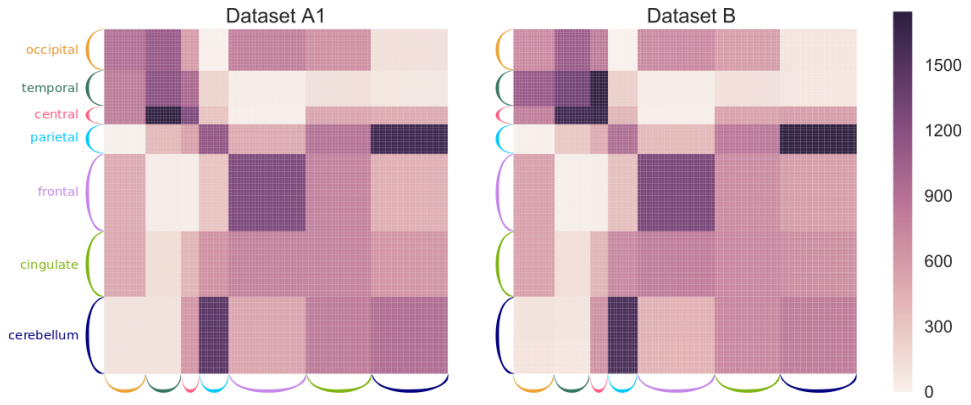

Figure 4: Correspondence of links at the level of subsystems. Number of links in each subsystem is represented by color for the two datasets A1 and B. The number of links for each subsystem is very similar between the two datasets.

with respect to both subjects and conditions using different planes of the data. In the high dimensional case different hyperplanes would be used, in practice different sets of links support the two classifications. Using MLR and EC, we achieved very high performance (accuracy >90%) for subject identification and perfect classification for the condition.

We then sought the smallest subsets of links that achieved the maximum performance of each classification using RFE (Figure 5B), as done before. Both support networks were again very sparse and distributed across the brain, as can be seen in their adjacency matrix (Figure 5C). More links are necessary to identify the subjects (57) than the behavioral conditions (13), indicating a higher complexity for the former.

Despite a (small) overlap of links between two networks, links relevant for subjects' identity and behavioral condition belong to almost disjoint sets. In order to prove this point we used RFE to rank the links according to their contribution to the classification, as we did before for datasets A1 and B. We computed the number of common links for the subject and condition identifications, which fell within the expected values for the null hypothesis (Figure 5D). Thus the overlap at the level of individual links is not any larger than that expected by chance.

Similar to Datasets A1 and B, subject identification of Dataset C largely concerns the frontal and cingulate systems. Condition identification is also supported by occipital and temporal cortices, which are expected to have the strongest activity modulations during movie viewing. The top panels in Figures 6A and B represent the two support networks such that the directed nature of links can be appreciated. Apart from two small components, the subject network appears almost fully connected with several central nodes (hubs, indicated by their large size), located in frontal and cingulate regions. In comparison, the condition network is segregated into small isolated components. The bottom plots in Figure 5 show the lateralization of the support links, stressing the asymmetries between the two hemispheres: most of the important links are ipsilateral (i.e., within the same hemisphere) and many belong to the left hemisphere for the subject network, whereas they are mainly contralateral for the condition network.

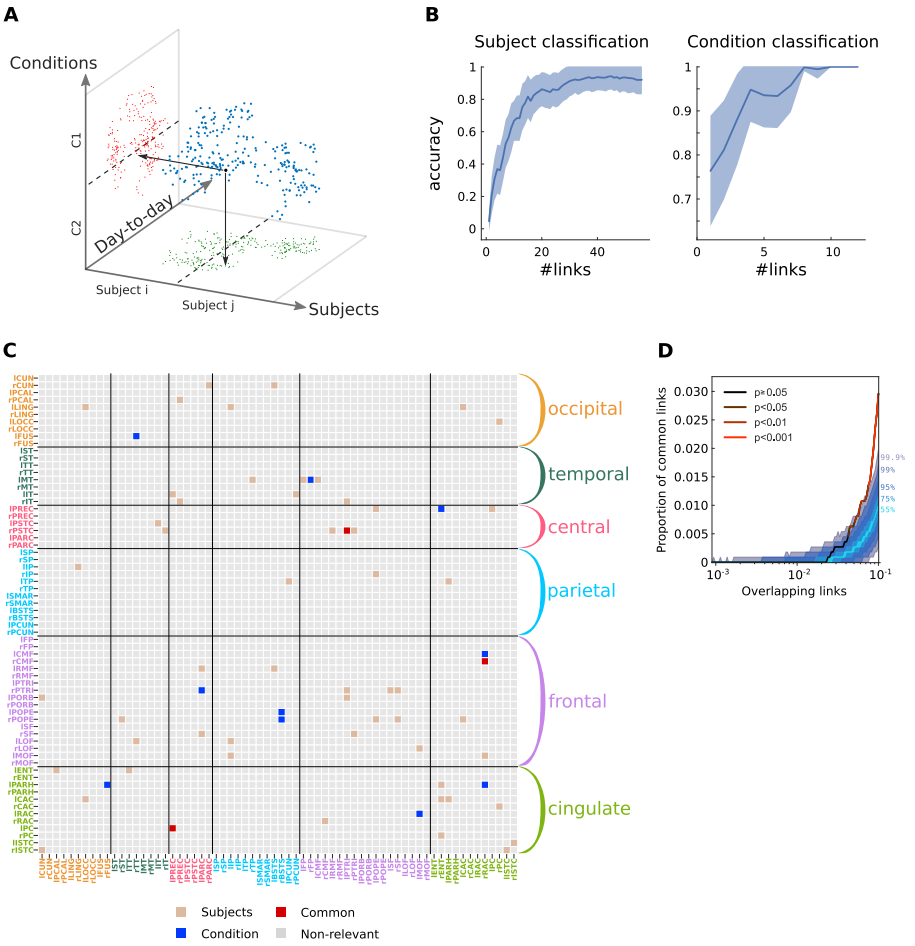

Figure 5: Twofold discrimination between subjects and conditions using EC. A) Idealized scheme of the twofold classification where each session (blue dots) is projected onto two planes, one for subjects (green) and one for conditions (red). In each plane, classification can be performed efficiently. Depending on the orthogonality of the subspaces, the two signatures have more or less overlap. B) Performance of the classification for 19 subjects and 2 conditions using Dataset C as a function of number of links. C) Signatures of the most discriminative EC links for the twofold classification: 54 links for subject classification in brown, 10 for condition classification in blue, 3 common links in red. D) Proportion of common links between the subject and condition signatures as a function of selected links. Color coding is the same as in Figure 4B.

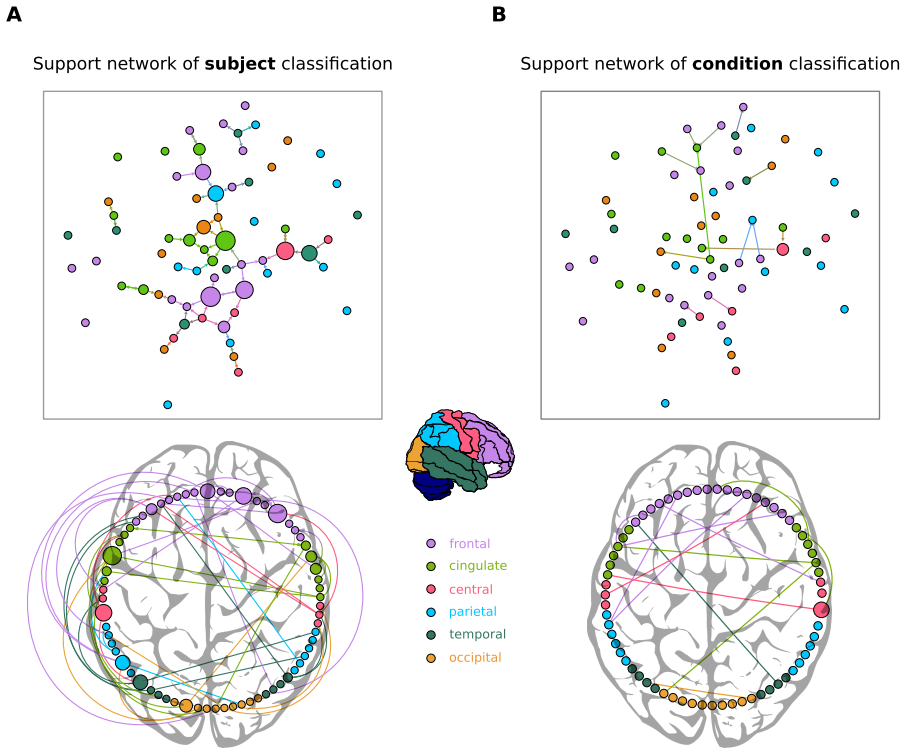

Figure 6: Support networks of subject and condition classification. A) The top graph plot represents the 57 most discriminative EC links supporting the classification of subjects (same as in Figure 5C). The size of each node represents its centrality in the extracted network (how much connected it is). The most central regions are located mainly in the frontal and cingulate cortices. The bottom circular plot shows the asymmetry and lateralization of the network, with more links located in the left hemisphere. Links that are inside the circle correspond to contralateral connections, while links outside the circle correspond to ipsilateral connections. B) Similar graph and circular plots as A for the 13 links supporting the classification between the two conditions.

3 Refinements

We presented a reliable method to classify simultaneously subjects and conditions and to extract the networks underlying these classifications. To the aim of applying this method to the study of different cognitive functions and to the clinical domain some refinements are needed.

The typical timespan of cognitive tasks is usually less than one minute (even few seconds for very stereotyped perceptual tasks). Our method is calibrated for fMRI sessions of few minutes. Even if experimental manipulations could provide tasks with longer time scales, an improvement of the method is necessary to allow the use of very short recording sessions. The main limitation of the method in this respect is related to the estimation of EC. Indeed the estimation procedure relies on the calculation of the correlation matrix. However when the number of ROIs is in the same order of the number of time points the correlation matrix will be very noisy (not to mention the case when there are more ROIs than time points, for which the correlation matrix becomes singular) and, as a consequence not reliable for the identification of subjects and conditions. Regularization is usually the solution to this type of problems, i.e. the pruning of the number of ROIs in this case. Another approach, the one that we are currently following, is that of a Bayesian estimation of the EC. The Bayesian setting allows the use of a prior probability on the EC that acts as a regularizer and allows a more robust estimation even when the number of time points is limited. The Bayesian approach allows also the estimation of the probability distribution of all parameters, through the so called Bayesian inference. This would allow in turn to test if a given connection exist or not thereby providing a more reliable estimation.

The clinical domain might also require some adjustments to the pipeline. We tested our method with behavioral conditions but the noise distribution in clinical conditions might be different. The use of PCA, maybe with a small number of principal components, should be evaluated in this case. In addition the use of non-linear classifiers could be beneficial and will be evaluated for each dataset.

References

[1] Noriaki Yahata, Kiyoto Kasai, and Mitsuo Kawato. Computational neuroscience approach to biomarkers and treatments for mental disorders. *Psychiatry and Clinical Neurosciences*, 71(4):215–237, 2017.

[2] Katrin Amunts, Michael J Hawrylycz, David C Van Essen, John D Van Horn, Noam Harel, J-B Poline, Federico De Martino, Jan G Bjaalie, Ghislaine Dehaene-Lambertz, Stanislas Dehaene, et al. Interoperable atlases of the human brain. *Neuroimage*, 99:525–532, 2014.

[3] B Biswal, FZ Yetkin, VM Haughton, and JS Hyde. Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. *Magn Reson Med*, 34:537–541, 1995.

[4] Dante Mantini, Uri Hasson, Viviana Betti, Mauro G Perrucci, Gian Luca Romani, Maurizio Corbetta, Guy A Orban, and Wim Vanduffel. Interspecies activity correlations reveal functional correspondence between monkey and human brain areas. *Nat Methods*, 9:277–282, 2012.

[5] Evan M Gordon, Timothy O Laumann, Adrian W Gilmore, Dillan J Newbold, Deanna J Greene, Jeffrey J Berg, Mario Ortega, Catherine Hoyt-Drazen, Caterina Gratton, Haoxin Sun, et al. Precision functional mapping of individual human brains. *Neuron*, 95(4):791–807, 2017.

[6] Mehdi Rahim, Bertrand Thirion, Danilo Bzdok, Irene Buvat, and Gaël Varoquaux. Joint prediction of multiple scores captures better individual traits from brain images. *NeuroImage*, 158:145–154, 2017.

[7] Sophie Kurth, Evelyne Moyse, Mohamed A Bahri, Eric Salmon, and Christine Bastin. Recognition of personally familiar faces and functional connectivity in alzheimer's disease. *Cortex*, 67:59–73, 2015.

[8] Mario Pannunzi, Rikkert Hindriks, Ruggero G Bettinardi, Elisabeth Wenger, Nina Lisofsky, Johan Martensson, Oisin Butler, Elisa Filevich, Maxi Becker, Martyna Lochstet, et al. Resting-state fmri correlations: from link-wise unreliability to whole brain stability. *Neuroimage*, 157:250–262, 2017.

[9] Sophia Mueller, Danhong Wang, Michael D Fox, Ruiqi Pan, Jie Lu, Kuncheng Li, Wei Sun, Randy L Buckner, and Hesheng Liu. Reliability correction for functional connectivity: Theory and implementation. *Human brain mapping*, 36(11):4664–4680, 2015.

[10] Elisa Filevich, Nina Lisofsky, Maxi Becker, Oisin Butler, Martyna Lochstet, Johan Martensson, Elisabeth Wenger, Ulman Lindenberger, and Simone Kühn. Day2day: investigating daily variability of magnetic resonance imaging measures over half a year. *BMC Neurosci*, 18:65, 2017.

[11] Emily S Finn, Xilin Shen, Dustin Scheinost, Monica D Rosenberg, Jessica Huang, Marvin M Chun, Xenophon Papademetris, and R Todd Constable. Functional connectome fingerprinting: identifying individuals using patterns of brain connectivity. *Nature neuroscience*, 18(11):1664–1671, 2015.

[12] Tobias Kaufmann, Dag Alnæs, Nhat Trung Doan, Christine Lycke Brandt, Ole A Andreassen, and Lars T Westlye. Delayed stabilization and individualization in connectome development are related to psychiatric disorders. *Nature neuroscience*, 20(4):513–515, 2017.

[13] Choong-Wan Woo, Luke J Chang, Martin A Lindquist, and Tor D Wager. Building better biomarkers: brain models in translational neuroimaging. *Nature neuroscience*, 20(3):365–377, 2017.

[14] Gustavo Deco and Morten L Kringelbach. Great expectations: using whole-brain computational connectomics for understanding neuropsychiatric disorders. *Neuron*, 84(5):892–905, 2014.

[15] Jesse Rissman and Anthony D Wagner. Distributed representations in memory: insights from functional brain imaging. *Annual review of psychology*, 63:101–128, 2012.

[16] Vince D Calhoun, Stephan M Lawrie, Janaina Mourao-Miranda, and Klaas E Stephan. Prediction of individual differences from neuroimaging data. *NeuroImage*, 145(Pt B):135–136, 2017.

[17] Matthew J Hoptman, Xi-Nian Zuo, Debra D'Angelo, Cristina J Mauro, Pamela D Butler, Michael P Milham, and Daniel C Javitt. Decreased interhemispheric coordination in schizophrenia: a resting state fmri study. *Schizophrenia research*, 141(1):1–7, 2012.

[18] Joerg F Hipp, Andreas K Engel, and Markus Siegel. Oscillatory synchronization in large-scale cortical networks predicts perception. *Neuron*, 69:387–396, 2011.

[19] Viviana Betti, Stefania Della Penna, Francesco de Pasquale, Dante Mantini, Laura Marzetti, Gian Luca Romani, and Maurizio Corbetta. Natural scenes viewing alters the dynamics of functional connectivity in the human brain. *Neuron*, 79:782–797, 2013.

[20] K Friston. Functional and effective connectivity: A review. *Brain Connect*, 1:8, 2011.

[21] Emily S Finn, Dustin Scheinost, Daniel M Finn, Xilin Shen, Xenophon Papademetris, and R Todd Constable. Can brain state be manipulated to emphasize individual differences in functional connectivity? *NeuroImage*, 2017.

[22] Oscar Miranda-Dominguez, Brian D Mills, Samuel D Carpenter, Kathleen A Grant, Christopher D Kroenke, Joel T Nigg, and Damien A Fair. Connectotyping: model based fingerprinting of the functional connectome. *PLoS One*, 9(11):e111048, 2014.

[23] Vicente Pallares, Andrea Insabato, Ana Sanjuan, Simone Kuehn, Dante Mantini, Gustavo Deco, and Matthieu Gilson. Subject- and behavior-specific signatures extracted from fmri data using whole-brain effective connectivity. *bioRxiv*, 2017.

[24] Matthieu Gilson, Gustavo Deco, Karl Friston, Patric Hagmann, Dante Mantini, Viviana Betti, Gian Luca Romani, and Maurizio Corbetta. Effective connectivity inferred from fmri transition dynamics during movie viewing points to a balanced reconfiguration of cortical interactions. *NeuroImage*, 2017.

[25] Matthieu Gilson, Ruben Moreno-Bote, Adrián Ponce-Alvarez, Petra Ritter, and Gustavo Deco. Estimation of directed effective connectivity from fMRI functional connectivity hints at asymmetries of cortical connectome. *PLoS Comput Biol*, 12:e1004762, 2016.

[26] Tamara Vanderwal, Jeffrey Eilbott, Emily S Finn, R Cameron Craddock, Adam Turnbull, and F Xavier Castellanos. Individual differences in functional connectivity during naturalistic viewing conditions. *Neuroimage*, 157:521–530, 2017.